

SPECIFICATION

TITLE OF THE INVENTION

5 CONTROL NETWORK WITH MATRIX ARCHITECTURE

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation application of U.S. Application Serial No.
09/442,368, filed on November 17, 1999, presently allowed. The foregoing application is hereby
10 incorporated by reference as if set forth fully herein.

BACKGROUND OF THE INVENTION

1) Field of the Invention

15 [0002] The field of the invention pertains to methods and apparatus for implementing a
control network and, more particularly, to a control network architecture and methods for
connecting nodes in a control network.

2) Background

20 [0003] Automated control systems are commonly used in a number of manufacturing,
transportation, and other applications, and are particularly useful to control machinery, sensors,
electronics, and other system components. For example, manufacturing or vehicular systems
may be outfitted with a variety of sensors and automated electrical and/or mechanical parts that
require enablement or activation when needed to perform their predefined functions. Such
25 systems commonly require that functions or procedures be carried out in a prescribed order or
with a level of responsiveness that precludes sole reliance on manual control. Also, such systems
may employ sensors or other components that require continuous or periodic monitoring and
therefore lend themselves to automated control.

[0004] As the tasks performed by machinery have grown in number and complexity, a need has arisen for ways to exercise control over the various components of a system rapidly, efficiently and reliably. The sheer number of system components to be monitored, enabled, disabled, activated, deactivated, adjusted or otherwise controlled can lead to difficulties in designing and implementing a suitable control system. As the number of system components to be controlled is increased, not only is the operation of the control system made more complicated, but also the wiring and inter-connections of the control system are likewise more elaborate. In addition, greater reliance on automated control has resulted in larger potential consequences if the automated control system fails.

[0005] Traditionally, control systems in certain applications, such as transit vehicles and railcars, have relied upon relay-based control technology. In such systems, relays and switches are slaved to a logic circuit that serves to switch signal connections. This approach requires a large number of relays and a substantial amount of wiring throughout the vehicle. In some instances distributed processors or logic circuits may be used for subsystems such as the door, but these processors or logic circuits often take up significant space and can be costly to maintain.

[0006] A substantial improvement has recently been made in the field of control systems. An improved network control system recently developed uses a dual-bus architecture along with distributed controllers. In this improved network control system, a primary bus forms a high-speed, bi-directional communication link interconnecting a main data bus controller with distributed slave modules, one of which acts as a second data bus controller connected to a secondary, low-speed data bus. The slave modules are generally connected to various input/output ports. The second data bus controller can be connected to second-tier slave modules over the secondary, low-speed data bus. The main data bus controller, secondary data bus controller, first-tier slave modules, second-tier slave modules, input/output ports and other system components collectively form a hierarchical system wherein the main data bus controller supervises the first-tier slave modules, including the second data bus controller, the second data bus controller supervises the second-tier slave modules, and the first-tier slave modules and second-tier slave modules supervise their assigned input/output functions.

[0007] While the dual-bus control network as described above has many advantages, there are also ways in which it could be improved further. The dual-bus control network architecture as currently known in the art generally relies on a single top-level main data bus controller. If the main data bus controller fails, system performance will be adversely impacted.

5 Also, the possibility of a short circuit occurring, particularly over a region of the bus, is a constant danger. In addition to disrupting communication signals among the components accessing the bus, a short circuit can be difficult to trace and cause substantial disruption of system service while maintenance personnel attempt to locate the short circuit. Furthermore, while the dual-bus network control architecture reduces wiring needed in a vehicle or other
10 automated system, simplification of wiring connections would lead to greater ease of implementation and maintenance.

[0008] Accordingly, it would be advantageous to provide a network control system that has a means for recovering from a failure in a main data bus controller or otherwise mitigating the effects such a failure. It would further be advantageous to provide a network control system
15 that reduces the impact of a short circuit and enables rapid identification of the location of a short circuit by maintenance personnel. It would further be advantageous to provide a distributed network control system with simplified wiring and connections.

[0009] In addition, as hierarchical control networks grow in size and complexity, it becomes increasingly difficult for the master node or controller to communicate with nodes
20 located remotely, separated often by a number of intervening, mid-level nodes or other intermediaries. Accordingly, it would be advantageous to provide a hierarchical control network which facilitates communication between the master node and lower level nodes.

SUMMARY OF THE INVENTION

25 **[0010]** According to one aspect of the invention, a control network is provided in which communication is facilitated between a master node and lower level nodes. In a preferred embodiment as described herein, a multiple-bus hierarchical control network includes a first-tier common bus and a plurality of lower-tier common buses. A first-tier master node controls a

plurality of first-tier slave nodes using the first-tier common bus for communication. Each of the first-tier slave nodes may be connected to a separate second-tier common bus, and each operates as a respective second-tier master node for a plurality of second-tier slave nodes connected to the particular second-tier common bus associated with the first-tier slave/second-tier master node.

5 Likewise, each of the second-tier slave nodes may be connected to a separate third-tier common bus, and each would then operate as a respective third-tier master node for a plurality of third-tier slave nodes connected to the particular third-tier common bus associated with the second-tier slave/third-tier master node. A supervisory network is connected to one or more of the lower-tier buses as well as to the first-tier common bus, so that activity on the lower-tier buses can be
10 monitored by the first-tier master node.

[0011] In a preferred embodiment, the supervisory network comprises a first-tier supervisory master node connected to a common supervisory data bus. A plurality of supervisory slave nodes are connected to the supervisory data bus. Each of the supervisory slave nodes is also connected to one of the lower-tier buses, and each supervisory slave node monitors the
15 communications occurring on the particular lower-tier bus. Information monitored by the supervisory slave nodes is communicated over the supervisory data bus to the first-tier supervisory master node either as needed, or on a regular basis, and can be relayed to the first-tier master node by the first-tier supervisory master node over the first-tier common bus.

[0012] If the control network is large and/or has many levels, then the supervisory
20 network can be expanded by providing additional lower-tier supervisory levels, including lower-tier supervisory buses, configured with a control structure similar to the hierarchical control network absent the supervisory network. The resulting preferred network structure may be viewed as having a “matrix” architecture due to the intersection of control paths emanating downward from a single master node with the supervisory paths emanating across the control
25 network from a single supervisory master node.

[0013] In a preferred embodiment, communication over a common bus is carried out using time division multiplexing. A time slot may be set aside for communication between the master node and the supervisory node connected to the common bus.

[0014] In another embodiment, a multi-master data bus supervisory network is connected to a hierarchical master-slave control network, wherein the multi-master data bus supervisory network monitors the activity on the various buses of the hierarchical master-slave control network. Activity on the various buses of the hierarchical master-slave control network is communicated to the upper tier master node(s) by relaying the necessary information over a multi-master data bus of the multi-master data bus supervisory network, which is generally much faster than attempting to relay the same information upwards tier-by-tier in the hierarchy of the hierarchical master-slave control network.

[0015] Various other embodiments as described herein provide for redundant backup control of the active master node on a common data bus. In a preferred embodiment, should a failure of the first-tier master node occur, any of the first-tier slave nodes connected to the first common bus can take over the first-tier master node, doing so according to their programmed priority. Should a failure of the second-tier master node occur, any of the second-tier slave nodes connected to a second common bus can take over the second-tier master node, doing so according to their programmed priority. Redundant master control is thereby provided for both the first tier and second tier in the hierarchical control network. Similar redundant backup control may be provided for any other common buses provided in the system organized in a master-slave hierarchy.

[0016] A preferred node comprises two separate transceivers, an uplink transceiver for receiving control information, and a downlink transceiver for sending out control information. Each node therefore has the capability of performing either in a master mode or a slave mode, or in both modes simultaneously.

[0017] In a preferred embodiment of the invention, a master node serves as a controller for a multiplicity of slave nodes. The master node polls the slave nodes periodically. Each of the slave nodes comprises a failure mode detector whereby, if a slave node fails to receive a message from the master node within a certain fixed period of time, then the slave node takes over control for the master node.

[0018] The invention provides in one aspect an automatic redundant backup master control for a master control node in a distributed, intelligent control network.

[0019] Further variations and embodiments are also disclosed herein, and are described hereinafter and/or depicted in the figures.

BRIEF DESCRIPTION OF THE DRAWINGS

5

[0020] Fig. 1 is a diagram of a distributed control network with two data buses as known in the prior art.

[0021] Fig. 2 is another diagram of a distributed control network having a two data buses each configured in a loop configuration as known in the prior art.

10 **[0022]** Fig. 3 is a circuit block diagram of a node that may be employed in the distributed control network of Fig. 1 or Fig. 2.

[0023] Fig. 4 is a diagram showing a physical encasement of the node shown in Fig. 3.

[0024] Fig. 5 is a block diagram of a preferred control network architecture in accordance with one or more aspects of the present invention.

15 **[0025]** Fig. 6 is a block diagram of a preferred node within the control network architecture shown in Fig. 5.

[0026] Fig. 7 is a diagram of a hierarchical control network in accordance with one embodiment of the present invention having multiple second-tier buses.

20 **[0027]** Fig. 8 is a diagram of a hierarchical control network in accordance with another embodiment of the present invention having a third-tier bus.

[0028] Fig. 9 is a functional diagram of a multi-bus control network illustrating one example of bus architectural layout and node functionality according to one embodiment of the invention.

[0029] Fig. 10 is a diagram of a hierarchical control network extrapolated to N tiers.

25 **[0030]** Fig. 11 is a diagram of a hierarchical control network with a matrix architecture according to an embodiment of the invention, so as to provide supervision of mid- and lower-tier levels of the control network.

[0031] Fig. 12 is a diagram of another embodiment of a hierarchical control network having a matrix architecture in accordance with various aspects of the invention.

[0032] Figs. 13A and 13B are timing diagrams for communication in a hierarchical control network having a matrix architecture, such as illustrated in Fig. 11 or 12.

[0033] Fig. 14 is a conceptual diagram of a multi-master data bus supervisory network connected to a hierarchical master-slave control network for the purpose of monitoring and/or supervising the various buses at different level tiers of the hierarchical master-slave control network.

[0034] Fig. 15 is a diagram of a hierarchical control network having a matrix architecture and utilizing one or more multi-master buses.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT(S)

[0035] This application is generally related to U.S. Patent 5,907,486 entitled "Wiring Method and Apparatus for Distributed Control Network," U.S. Patent Application Ser. No. 08/854,160 filed in the name of inventor Jeffrey Ying, entitled "Backup Control Mechanism in a Distributed Control Network," U.S. Patent Application Ser. No. 08/853,893 filed in the name of inventors Jeffrey Ying and Michael Kuang, entitled "Fault Isolation and Recovery In A Distributed Control Network," and U.S. Patent Application Ser. No. 08/853,989 filed in the name of inventor Jeffrey Ying, entitled "Multi-Tier Architecture for Control Network," all of which foregoing are hereby incorporated by reference as if set forth fully herein.

[0036] Figure 1 is a block diagram showing the interconnection of nodes in a particular type of control network 101 as known in the art. The control network 101 comprises a main data bus controller 103 which is connected over a main data bus 104 to a plurality of first-tier slave nodes 109 and 123. One first-tier slave node 123 connected to the main data bus 104 also functions as a second data bus controller, and is connected to a second data bus 113. The second data bus controller 123 is connected over the second data bus 113 to a plurality of second-tier slave nodes 130. The main data bus 104 forms a high-speed, bi-directional communication link between the main data bus controller 103 and the first-tier slave nodes 109 and 123, and the second data bus 113 forms a low-speed, bi-directional communication link between the second data bus controller 123 and the second-tier slave nodes 130.

[0037] The nature of the slave nodes 109, 123 and 130 depends in part on the control application for which they are deployed. In a transit vehicle or railcar, for example, the master data bus controller 103 and the slave nodes 109, 123 and 130 may each be assigned to control a particular section of the vehicle or railcar, or may be assigned to control particular input and output functions. For each slave node 109, 123 and 130 in Fig. 1, various control signals are shown connected to the nodes such as to illustrate one exemplary arrangement of control functionality.

[0038] In operation, the main controller 103 communicates with the first-tier slave nodes 109 and 123 using the main data bus 104 as a high speed bi-direction link. An exemplary baud rate for communications over the main data bus 104 is 256k. The main data bus controller 103 is generally responsible for delegating control commands to the first-tier slave nodes 109 and 123, and for responding to status information and events communicated to the main data bus controller 103 over the main data bus 104. Each of the first-tier slave nodes 109 and 123 receives commands from the main data bus controller 103, and issues appropriate commands over their respective control lines. In a similar manner, the second data bus controller 123 communicates with the second-tier slave nodes 130 using the second data bus 113 as a low speed bi-direction link (having a baud rate of, e.g., 9.6k), and instructs the second-tier slave nodes 130 to carry out certain control functions, or responds to status messages or events relayed to the second data bus controller 123 from the second-tier slave nodes 130.

[0039] Figure 2 is a diagram showing the layout or architecture of the Fig. 1 control network. The control network 201 shown in Fig. 2 comprises a main data bus controller 203 which is connected to a main data bus 204. The main data bus 204 is physically connected to a plurality of first-tier slave nodes 209 and 223. As explained with respect to the control network 101 shown in the Fig. 1, one of the first-tier slave nodes 223 also functions as a second data bus controller 223, and is connected over a second data bus 213 to a plurality of second-tier slave nodes 230. The main data bus 204 is configured in a loop such that it passes through each of the first-tier slave nodes 209 and 230 and returns to rejoin the main data bus controller 203. In this way, should the wires of the main bus 204 become severed, the main data bus controller 203 will still be connected to the first-tier slave nodes 209 and 223 and will not necessarily lose control

over the system. Similarly, the second data bus 213 is configured in a loop such that it passes through each of the second-tier slave nodes 230 and returns to rejoin the second data bus controller 223, thereby providing an architecture resilient to potential severing of the wires of the second data bus 113. Each of the main data bus controller 203, first-tier slave nodes 209 and 223, and second-tier slave nodes 230 may be connected to a plurality of control signals for performing control or sensor functions, or various other input and output functions as necessary for the particular control application.

[0040] The control network 201 shown in Fig. 2 thus utilizes a dual-bus architecture to perform control functions. Because of the hierarchical architecture of the control system 201, relatively low baud rates on the second data bus 213 can be tolerated, leading to reduced system size, cost and complexity over traditional non-hierarchical, relay-based systems. The slower speed on the secondary data bus 213 also reduces the system's susceptibility to electromagnetic interference, a potential problem in certain control system environments (such as railcars).

[0041] Each node, whether master data bus controller 203, first-tier slave node 209 or 223, or second-tier slave node 230, includes means for performing computations necessary for its functionality, and is configured with components such as a central processing unit (CPU) and memory. Figure 3 is a more detailed block diagram of a node 301 (such as the master data bus controller 203, a first-tier slave node 209 or 223, or a second-tier slave node 230) that may be employed in the control network of Fig. 2. The node 301 comprises a CPU 315 connected to a power control block 317 and a transceiver 305. The node 301 is also connected to power signal lines 316, which connect to the power control block 317. The node 301 may communicate over communication signal lines 304, which are connected to the transceiver 305. An electrical erasable programmable read-only memory (EEPROM) 306 stores programming information utilized by the CPU 315 for carrying out certain programmable functions. The CPU 315 has access to a random access memory (RAM) (not shown) and read-only memory (ROM) (not shown) as needed for the particular application.

[0042] The CPU 315 is connected to a keyboard and display interface block 320. The keyboard and display interface block 320 is connected to status LEDs 307, relays 321, and LED display 311 and a keypad 331. The node 301 is thereby can accept manual inputs (e.g., from the

keypad 331) or receive sensor inputs (e.g., over relays 321), and can display operational status using status LEDs 301 or LCD display 311.

[0043] The node 301 further comprises a network controller 322 which preferably comprises a second CPU. The network controller 322 is connected to a second transceiver 323 which is connected to a second pair of communication signal lines 314. The network controller also outputs power signal lines 336.

[0044] In operation, node 301 may communicate over two different data buses using transceivers 305 and 323. Thus, node 301 may communicate over a first data bus (such as data bus 204 shown in Fig. 1) by receiving and transmitting signals over communication signal lines 314 using transceiver 323, under control of the network controller 322. The node 301 may communicate over a second data bus (such as data bus 213 shown in Fig. 2) by transmitting and receiving signals over communication signal lines 304 using transceiver 305, under control of CPU 315. The CPU 315 and network controller 322 may transfer information back and forth using a shared memory (not shown). The node 301 may serve as both a "slave" unit with respect to the first data bus 204 and a "master" unit with respect to the second data bus 213. By interconnecting a plurality of nodes 301 in an appropriate configuration, a hierarchical control network with two data buses (as shown in Fig. 2) may be established.

[0045] Each node 301 such as shown in Fig. 3 is housed in a rugged, potted case made of a suitable lightweight material such as aluminum that provides environmental protection and allows for heat dissipation. Figure 4 is a diagram showing an exemplary physical casing 401 of a module or node 301 such as shown in Fig. 3. The casing 401 can be quite small; in the example of Fig. 4, the casing 401 measures approximately 2.1" by 3.75", and is 0.825" in thickness.

[0046] A problem that can occur in operation of a control network such as shown in Fig. 2 is that if the master data bus controller 203 fails then operation of the entire system could be jeopardized. A possible solution would be to provide a redundant master data bus controller that has the same functionality as the primary master data bus controller 203 in all respects. Upon detecting a failure of the primary master data bus controller 203, the backup master data bus controller could shut down the primary master data bus controller 203 and take over control of the network.

[0047] While having such a separate, redundant master data bus controller for backup purposes may provide a solution where the primary master data bus controller 203 fails, it falls short of being a complete solution. As an entirely separate controller having complete functional and hardware redundancy of the primary master data bus controller 203, incorporation of the backup master data bus controller effectively doubles the cost of implementing the master data bus controller 203. Also, another drawback is that if both the master data bus controller 203 the backup master data bus controller fail, then operation of the entire system would be jeopardized and operation could come to complete halt.

[0048] In addition to the possibility of the master data bus controller 203 failing, the second data bus controller 223 could also be subject to failure. While a redundant second data bus controller for backup purposes could be provided, the cost of implementing the second data bus controller would be essentially doubled, and the system is still subject to potentially complete failure should the second data bus controller also fail. Moreover, adding redundant data bus controllers could complicate the wiring of the system.

[0049] A preferred embodiment of the invention overcomes one or more of the above problems by providing redundant backup control for the master data bus controller 203 or other type of master node, the second data bus controller 223 or similar types of nodes, and, if further nested control levels exist (as described, for example, in later embodiments herein), other sub-controllers for those control levels.

[0050] Figure 5 is a block diagram of a preferred embodiment of a control network 501 having redundant backup control capability for a master node at each bus level of the control network 501. Hereinafter, the node acting as the master bus controller for a particular bus will be referred to as the "master node" for that particular bus, and all the other nodes on that bus will be referred to as "slave nodes" for that particular bus. In the control network shown in Fig. 5, a master node 503 and a plurality of first-tier slave nodes 523 are connected to a main data bus 504. In a preferred embodiment of the invention, each of the slave nodes 523 is configured or can be configured to control a secondary data bus. For example, the first-tier slave node 523c is shown connected to a secondary data bus 523 in the control network 501. The first-tier slave node 523c functions as a second-tier master node with respect to second-tier slave nodes 533

connected to the secondary data bus 513. Others of the first-tier slave nodes 523 can also serve as second-tier master nodes and be connected to different secondary buses having additional second-tier slave nodes. A multi-level or multi-tiered hierarchical control network is thereby established.

5 **[0051]** Each of the master node 503, first-tier slave nodes 523, second-tier slave nodes 533, and other lower-level slave nodes (not shown in Fig. 5) are referred to hereinafter generically as "nodes" and are designated as nodes 530 in Fig. 5. In one aspect of a preferred embodiment as shown in Fig. 5, each of the nodes 530 has substantially the same hardware configuration and can therefore function as either a master node or a slave node, depending upon
10 how the control network 501 is configured. Each data bus, along with the nodes attached to it, are generally referred to as a cell, and the master node connected to the data bus is referred to as a "cell controller" for that particular cell. As explained in more detail hereinafter, each node 530 configured as a master node transmits and receives messages over the data bus for the cell it controls. Each node 530 configured as a slave node remains in a listen mode, receiving but not
15 transmitting messages over that data bus, unless specifically requested to transmit information over the data bus by the master node. Any number of the slave nodes can, even though operating as a slave node with respect to an upper tier, be simultaneously operating as a master node with respect to other lower-tier slave nodes at a different cell sub-level.

[0052] A preferred embodiment of the invention, as noted, comprises a mechanism for
20 redundant backup control of any node functioning as a master node at any level or sub-level of the control network 501. As generally described, in operation of a preferred embodiment of the invention the slave nodes connected to a particular data bus monitor the data bus while in a listen mode and await periodic signals from the master node for that data bus. Upon a failure to receive a signal from a master node within an expected time, the slave nodes connected to that data bus
25 begin a wait period (which is preferably a different wait period for each slave node connected to the data bus). When the wait period elapses, the slave node determines that a failure in the master node for the particular data bus has occurred, and takes steps to take over the functionality of the master node. Each of the slave nodes is programmed with a different wait period, so that there is no contention for replacing the master node when a master node failure has occurred. In

one aspect, backup control of each master node is prioritized, such that there is a specific order in which the slave nodes can potentially take over control of the master node functionality when a failure has occurred.

[0053] In more detail, again with reference to Fig. 5, one of the nodes 530 attached to the main data bus 504 is configured as a master node 503. The other nodes 530 attached to the main data bus 504 (in this example numbering four such nodes 530) are configured as first-tier slave nodes 523, meaning that they receive but do not transmit master-control signals over the main data bus 504. The first-tier slave nodes 523 may, however, from time to time send responsive signals or status signals over the main data bus 504.

[0054] In a preferred embodiment, each of the first-tier slave nodes 523 may be configured as a second-tier master node controlling a secondary bus. One such example is shown in Fig. 5, wherein first-tier slave node 523c is connected to a secondary data bus 513. A plurality of other nodes 530 are also attached to the secondary bus data 513, and serve as second-tier slave nodes 533. There are three such second-tier slave nodes 533 in the example shown in Fig. 5.

With respect to the secondary data bus 513, the first-tier slave/second-tier master node 523c transmits master-control signals to the second-tier slave nodes 533. The second-tier slave nodes 533 ordinarily operate only in a listen mode, but from time to time may send responsive messages or status messages to the second-tier master node 523c. The other first-tier slave nodes 523a, 523b and 523d may similarly be connected as second-tier master nodes (i.e., cell controllers) each controlling its own secondary bus or cell.

[0055] While the control network 501 shown in Fig. 5 has four first-tier slave nodes 523 and three second-tier slave nodes 533, the number of first-tier slave nodes 523 and second-tier slave nodes 533 is limited only by the ability of the master node to communicate with the slave nodes over the particular data bus. There may be more slave nodes or fewer slave nodes on each bus than shown in the control network 501. In a preferred embodiment, there are no more than eight such cell controllers, although more than eight may be used so long as processing capacity and speed permit.

[0056] In addition, further levels of control nesting beyond two data buses may also be provided, using a similar approach to the two data bus method. Thus, for example, one or more

of the second-tier slave nodes 533 may be configured as a third-tier master node controlling its own tertiary or third-tier data bus. While Fig. 5 only shows two nested control levels, the same control concepts would apply to a control network architecture having additional nested control levels. Examples of control networks having more than two data buses are depicted in Figs. 7, 8 and 9 and described in more detail hereinafter.

[0057] In a preferred embodiment, communication over the main data bus 504 and the secondary data bus 513 (or buses, if appropriate) is time-multiplexed such that only one node 530 is transmitting over a particular data bus at a given time. Usually, each transmitted message will be targeted for a specific destination node 530, which may be specified by address bits in the transmitted message. However, in some embodiments broadcast messages may also be used targeted to multiple nodes 530.

[0058] Responsibilities for tasks, or groups of tasks, may be assigned to specific nodes 530. For example, each of the first-tier slave nodes 223 may be assigned a distinct sphere of responsibility. Similarly, each of the second-tier slave nodes 533 may be assigned a distinct sphere of responsibility. Examples of tasks that may be assigned to different nodes 530 are described for an exemplary control network later herein, with respect to Fig. 9.

[0059] Each of the nodes 530 preferably comprises an uplink transceiver 507, a downlink transceiver 508, and a switch 509. Each of the nodes 530 receives signals over its downlink transceiver 508. Over the main data bus 504, the first-tier master node 503 transmits master-control signals to each of the first-tier slave nodes 523. From time to time, according to the programmed control protocol, the first-tier slave nodes 523 respond to the master-control signals, or otherwise send status messages to the first-tier master node 503 when events occur specific to that first-tier slave node 523. Otherwise, the first-tier slave nodes 523 do not ordinarily communicate with each other.

[0060] In a similar manner, over each secondary data bus (such as secondary data bus 513), the second-tier master node 523 (for example, first-tier slave/second-tier master node 523c in Fig. 5) transmits master-control signals to each of the second-tier slave nodes 533 connected to the same secondary data bus. From time to time, according to the programmed control protocol, the second-tier slave nodes 533 respond to the master-control signals, or otherwise send status

messages to the second-tier master node 523c when events occur specific to that second-tier slave node 533. Otherwise, the second-tier slave nodes 523 do not ordinarily communicate with each other.

[0061] Communication between nodes is preferably carried out using half-duplex time division multiplexing. In typical operation, the master node polls each of the slave nodes periodically. Each of the nodes is preferably provided with a unique node identification number or address that distinguishes it from all other nodes of the control network. The master node sends a control message to each slave unit in turn, using the node identification number or address to identify the intended destination. Each of the slave nodes receives the control message but only reacts if it recognizes its own node identification number or address in the control message. The slave node takes the actions requested by the control message received from the master node. Within a designated time period after receiving the control message, the slave node responds to the master node with an acknowledgment message. Each of the slave nodes are polled in turn so that the master node can keep track of events happening throughout the system.

[0062] A communication protocol is preferably established so as to avoid collisions on each of the data buses. A simple and effective communication protocol is one in which the master node for the particular data bus sends a control message to a particular slave node, which responds with an acknowledgment or status message within a predetermined amount of time before the master node contacts another slave node. Slave nodes generally do not initiate communication without being first polled by the master node. The master node may also send out a broadcast control message that is intended for receipt by more than one of the slave nodes. The broadcast control message can comprise a node identification number or address that instructs a single particular node to respond to the broadcast control message. Usually, the single node selected for response will be the most critical node requiring receipt of the broadcast control message.

[0063] Failure of the current master node (at any of the control levels) commonly results in the master node either failing to transmit, or else transmitting improper control information to the slave nodes over the data bus. According to a preferred redundant backup control protocol, the slave nodes periodically receive master-control messages from the master node and, in the

event that proper master-control messages fail to appear, initiate a failure mode response procedure.

[0064] Detection of and response to a failure mode condition may be explained in greater detail with reference to Fig. 6, which is a block diagram of a preferred embodiment depicting most of the main components of a node (such as any of nodes 530 shown in Fig. 5). Because failure mode detection and response is carried out by a node 530 operating as a slave node, the following discussion will assume that the node 603 shown in Fig. 6 is initially configured as a slave node. Further, for simplicity of explanation, it will be assumed that the node 603 shown in Fig. 6 is a first-tier slave/second-tier master node connected to a main bus and a secondary bus (such as first-tier slave/second-tier master node 523c connected to the main data bus 504 and secondary data bus 513 in Fig. 5), although the same node circuit configuration is preferably used for each of the nodes 530, regardless of control level, for ease of construction and flexibility purposes.

[0065] In the node block diagram of Fig. 6, a node 603 is shown connected to a first bus (e.g., main bus) 604. The node 603 comprises an uplink transceiver 611, a downlink transceiver 621, a CPU 612 connected to the uplink transceiver 611, and another CPU 622 connected to the downlink transceiver 621. Both CPUs 612, 622 are preferably connected to a dual-port RAM 618, and each CPU 612, 622 is connected to a ROM program store 614 and 624, respectively. The second CPU 622 is connected through an appropriate interface to I/O ports 654, which may comprise sensor inputs, control signal outputs, status LEDs, LCD display, keypad, or other types of external connections. It will be understood that the node 603 of Fig. 6 can have all the components and functionality of the node 301 shown in Fig. 3; however, in Fig. 6 only certain basic components needed for explaining the operation of the invention are depicted.

[0066] Each node 603 is preferably capable of both sending and receiving messages (e.g., control instructions). Typically, the uplink transceiver 611 operates in a “slave” mode whereby the node 603 receives control instructions using the uplink transceiver 611 and then responds thereto, and the downlink transceiver 621 operates in a “master” mode whereby the node 603 issues control instructions (e.g., polls slave nodes) and awaits a response from other nodes after sending such control instructions.

[0067] The downlink transceiver 621 of the node 603 is connected to a secondary data bus 652, to which is also connected a plurality of second-tier slave nodes 651 (assuming the node 603 is a first-tier slave/second-tier master node). The node 603 thereby functions as a first-tier slave node with respect to the main data bus 604, receiving with its uplink transceiver 611 first-tier master-control signals over the main bus 604 from a first-tier master node (such as master node 503 shown in Fig. 5), and also functions as a second-tier master node with respect to the secondary data bus 652, transmitting second-tier master-control signals with its downlink transceiver 634 to second-tier slave nodes 651.

[0068] The node 603 also comprises a pair of switches 635a, 635b connected between the downlink transceiver 621 and the signal lines 643a, 643b of the main data bus 604. In normal operation, the switches 635a, 635b remain open (unless the node 503 is also the first-tier master node, such as master node 503 shown in Fig. 5, in which case the switches 635a, 635b would be closed), and the downlink transceiver 621 is thereby isolated from the main data bus 604. However, when a first-tier master node failure condition is detected, switches 635a, 635b are closed, enabling the downlink transceiver 621 to take over for the first-tier master node. The downlink transceiver 621 would therefore function simultaneously as master node with respect to both the main data bus 604 and the secondary data bus 652.

[0069] In a preferred embodiment, detection of a master node failure condition on the main data bus 604 is accomplished using a timer mechanism, such as a hardware timer 613 accessible (either directly or indirectly) by the CPU 612 that is connected to the uplink transceiver 611. According to a preferred control protocol (assuming the node 603 is a first-tier slave/second-tier master node), the uplink transceiver 611 of node 603 receives first-tier master-control signals periodically from the first-tier master node (such as master node 503 in Fig. 5). The master-control signals may, for example, request status information from the node 603, or instruct the node 603 to carry out certain control or input/output functions. The node 603 ordinarily responds by carrying out the requested functions and/or sending an acknowledgment or status signal to the first-tier master control node using the uplink transceiver 611.

[0070] Timer 613 times out a wait period between master-control signals received from the first-tier master control node. In a preferred embodiment, each time the uplink transceiver

611 receives a master-control signal from the first-tier master node that is recognized as an appropriate master-control signal within the particular programmed control protocol (whether or not the master-control signal is directed to the particular node 603), the CPU 612 connected to the uplink transceiver 612 resets the timer 613. If the timer 613 ever times out, then CPU 612
5 responds by asserting a failure mode response procedure. The timing out of timer 613 may result in an interrupt to CPU 612 in order to inform the CPU 612 of the failure to receive master-control signals, or else the CPU 612 may periodically monitor the timer 613 and, when the CPU 612 notices that the timer 613 has timed out, assert a failure mode response procedure.

[0071] When a failure mode condition is detected, the CPU 612 sets a failure mode status
10 bit in a predetermined flag location within the dual-port RAM 618. The other CPU 622 periodically monitors the failure mode status bit in the dual-port RAM 618 and is thereby informed when a failure occurs. Alternatively, instead of the CPUs 612, 622 communicating through the dual-port RAM 618, timer 613 can directly inform CPU 622 when a failure to receive master-control signals has occurred (i.e., when timer 613 has timed out).

[0072] When the CPU 622 has been informed or otherwise determined that a failure
15 mode condition exists, and that the first-tier master node has presumably failed, the CPU 622 sends a signal over control line 633 to close switches 635a, 635b, thereby connecting the downlink transceiver 621 to the main bus 604. From that point on, the CPU 622 performs as the first-tier master node with respect to the main bus 604. The node 603 can continue to receive
20 information over the main data bus 604 using the uplink transceiver 611. Alternatively, the node 603 may thereafter perform all transmission and reception over both the main bus 604 and the secondary bus 652 using the downlink transceiver 621. When the failure mode is entered, the CPU 622 may be programmed so as to directly carry out the I/O port functions for which it previously received instructions from the first-tier master node, or the node 603 may send
25 master-control signals to its own uplink transceiver 611 and thereby continue to carry out the I/O port functions as it had previously been doing. In other words, the node 603 can give itself control instructions over the main data bus 604 so that it can continue to perform its previously assigned functions. If, after taking over for the first-tier master node, the node's downlink transceiver 611 should fail, the node 603 can still continue to perform its control functions when

the next slave node takes over control as the new first-tier master node (as later described herein), because its uplink transceiver 611 continues to function in a normal manner.

[0073] According to the above described technique, the node 603 thereby substitutes itself for the first-tier master node upon the detection of a first-tier master node failure as

5 indicated by the failure to receive the expected first-tier master-control signals. Should the node 603 fail, either before or after taking over control for the first-tier master node, the next first-tier slave node would take over and become the first-tier master node in a similar manner to that described above.

[0074] Referring again to Fig. 5, the order in which the first-tier slave nodes 523 take over for the first-tier master node 503 is dictated by the wait period timed out by the timer 613 of
10 the particular first-tier slave node 523. The timer 613 (see Fig. 6) for each first-tier slave node 523 is programmed or reset using a different time-out value. A first-tier slave node 523 only asserts a failure mode condition when its internal timer 613 reaches the particular timeout value programmed for that particular node 523.

15 **[0075]** While the programmed wait periods for the internal timer 613 in each first-tier slave node 523 can vary depending upon the control application, illustrative wait periods are programmed in ten millisecond increments. Thus, for example, first-tier slave node 523a could be programmed with a 10 millisecond wait period; the next first-tier slave node 523b could be programmed with a 20 millisecond wait period; the next first-tier slave node 523c could be
20 programmed with a 30 millisecond wait period; and the last first-tier slave node 523d could be programmed with a 40 millisecond wait period; and so on. First-tier slave node 523a would take over as the first-tier master node if 10 milliseconds elapses without it receiving any proper first-tier master-control signals; the next first-tier slave node 523b would take over as the first-tier master node if 20 milliseconds elapses without it receiving any proper first-tier master-control
25 signals; the next first-tier slave node 523c would take over as the first-tier master node if 30 milliseconds elapses without it receiving any proper first-tier master-control signals; and so on.

[0076] Use of 10 millisecond increments for the wait periods in the above example is considered merely illustrative, and the actual wait periods should be selected depending upon the time criticality of the control messages, and the number of messages that may be missed before a

high enough degree of certainty is established that the master node has failed. For example, if a slave node expects to observe a control-message signal on the data bus no later than every 5 milliseconds, then the slave node may be programmed to assert a failure mode condition after a wait period corresponding to the absence of a predefined number of messages--for example, twenty messages (i.e., 100 milliseconds). If critical aspects of the system requiring master node control need to be serviced in a shorter time period, then the wait period would have to be reduced to accommodate the time-sensitive components of the system.

[0077] The order in which the slave nodes take over for the master node need not be dictated by the relative position in the control loop of the slave node with respect to the master node, but rather may be dictated according to the programmed wait period in each slave node. Flexibility is thereby provided in the order of priority in which the slave nodes take over for the master node in the event of a failure event.

[0078] Accordingly, by use of the inventive techniques described herein, redundant backup for the first-tier master node 503 is provided. Such redundant backup control is provided without requiring additional physical nodes to be located within the control system, and without having to provide wiring for such additional physical nodes to the buses 504 or 513. The redundant backup for the master node 504 is also accomplished while resolving contention problems that might otherwise occur if each of the first-tier slave nodes 523 were programmed with the identical timeout period.

[0079] In a preferred embodiment, redundant backup control is provided in a similar manner for the secondary data bus 513, and each additional data bus that may be provided in the system (e.g., in systems such as shown in Figs. 7, 8 or 9). Thus, each of the second-tier slave nodes 533 is preferably configured with the circuitry shown for node 603 in Fig. 6, and each of the second-tier slave nodes 533 can therefore substitute itself for the first-tier slave/second-tier master node 523c if the first-tier slave/second-tier master node 523c fails.

[0080] If a particular node is operating as a master node for two buses as a result of a failure of the master node on a higher-tier bus, and the node operating as such fails, then it is possible that two different nodes will take over for the failed node, one node taking over on each bus. For example, supposing that first-tier slave/second-tier master node 523c has already taken

over as the first-tier master node due to a failure of the master node 503, and further suppose that first-tier slave/second-tier master node 523c too fails, then the next first-tier slave node 523d would take over as the first-tier master node with respect to the main data bus 504, but the first second-tier slave node 533a would take over as second-tier master node with respect to the
5 secondary data bus 513.

[0081] In the above manner, despite the failure of one or more nodes, substantial functionality of the control system as a whole can be maintained. A failed node is essentially discarded or bypassed to the extent possible so as to maintain the highest possible degree of continued operability. Furthermore, because certain parts of the system will continue operate
10 despite the failure of the master node, identification of the failed node by engineers or maintenance personnel should be simplified by being able to identify the inoperative portion of the system that has become isolated due to the failure.

[0082] In one aspect, separation of responsibility in each node 603 of master functions and slave functions between two different CPU's each operating with a different transceiver
15 allows the node 603 to potentially continue operating as either a master node or a slave node should one of the CPU's fail, providing that the failure does not disrupt both of the transceivers at the node 603.

[0083] In a preferred embodiment, the nodes 530 of Fig. 5 are wired using a single cable connecting all of the nodes 530 in a loop configuration. Details of such a wiring technique are
20 described in U.S. Patent 5,907,486 entitled "Wiring Method and Apparatus for Distributed Control Network," assigned to the assignee of the present invention, and previously incorporated herein by reference.

[0084] In a preferred embodiment, the nodes 530 of Fig. 5 are configured with fault isolation and recovery circuitry in the case of a short circuit or similar event. Details of such
25 fault isolation and recovery circuitry are described in copending U.S. Application Ser. No. 08/853,893 entitled "Fault Isolation and Recovery In A Distributed Control Network," previously incorporated herein by reference.

[0085] Figures 7, 8 and 9 depicts various embodiments having more than two data buses, so as to provide additional levels of control beyond that afforded by a dual-bus architecture.

Each of the nodes shown in Figs. 7, 8 and 9 is preferably configured to include the circuitry shown for preferred node 603 in Fig. 6. Figure 7 shows an example of a system architecture for a control network having three data buses 704, 714 and 724. A first-tier master node 703 and a plurality of first-tier slave nodes 712 are connected to the main data bus 704. One of the first-tier slave nodes 712, designated as A1 in Fig. 7, operates as a second-tier master node, and is connected to the second data bus 714 along with a plurality of second-tier slave nodes 722. Another of the first-tier slave nodes 712, designated as D1 in Fig. 7, operates as another second-tier master node, and is connected to the third data bus 724 along with another plurality of second-tier slave nodes 732. The other first-tier slave nodes 712, designated B1 and C1 in Fig. 7, could also be configured as master nodes of a second-tier bus. Figure 7 thereby provides a hierarchical control network 701 having two control levels or tiers, and three data buses.

[0086] Figure 8 shows an example of a system architecture for a control network having four buses 804, 814, 824 and 834. In a similar manner to Fig. 7, a first-tier master node 803 and a plurality of first-tier slave nodes 812 are connected to the main data bus 804. One of the first-tier slave nodes 812, designated as A1 in Fig. 8, operates as a second-tier master node, and is connected to the second data bus 814 along with a plurality of second-tier slave nodes 822. Another of the first-tier slave nodes 812, designated as D1 in Fig. 8, operates as another second-tier master node, and is connected to the third data bus 824 along with another plurality of second-tier slave nodes 832. One of the second-tier slave nodes 832 connected to the third data bus 824, denoted as A2' in Fig. 8, operates as a third-tier master node with respect to the fourth data bus 834, which is connected to a plurality of third-tier slave nodes 842. Figure 8 thereby provides a hierarchical control network 801 having three control levels or tiers, and four data buses.

[0087] It will be appreciated that, expanding the approach used in Figs. 7 and 8, additional control levels may be created by adding successive lower control tiers, or additional slave nodes at any particular tier may be configured as cell controllers to control additional localized data buses. A great deal of flexibility is thereby provided in establishing a hierarchical control structure suitable for many different control applications.

[0088] Figure 9 is a diagram showing, from a functional standpoint, an example of a particular control application having multiple data buses in accordance with the hierarchical control principles discussed herein. In Fig. 9, a control network 901 comprises a master node 904 which is connected to a plurality of slave nodes 923, 924, 925 and 926, each of which is assigned a particular sphere of responsibility within the control network. A main bus 903 forms a communication link between the master node 904 and the slave nodes 923, 924, 925 and 926.

[0089] Generally, the nature of the slave nodes 923, 924, 925 and 926 depends in part on the control application in which they are deployed. In the example of Fig. 9, the slave nodes 923, 924, 925 and 926 are deployed in a vehicle or railcar, and so the slave nodes 923, 924, 925 and 926 have functionality suited for such a control application. For example, the slave nodes include a slave node 923 operating as a rear section controller, a slave node 924 operating as a central section controller, a slave node 925 operating as a front section controller, and a slave node 926 operating as a panel controller. There may also be additional slave nodes if required.

[0090] Each of slave nodes 923, 924, 925 and 926 are considered first-tier slave nodes in the illustrative embodiment shown in Fig. 9. In the control network 901 of Fig. 9, two of the first-tier slave nodes 923, 924 also act as second-tier master nodes for additional data buses. Thus, first-tier slave node 923 operates as a second-tier master node with respect to a second data bus 913, and first-tier slave node 924 operates as a second-tier master node with respect to a third data bus 914. First-tier slave/second-tier master node 923 is connected to a plurality of second-tier slave nodes 931, 932, 933 and 934, which may each be assigned a sub-sphere of responsibility in the cell controlled by the rear section controller. The second-tier slave nodes may therefore include, for example, a slave node 931 operating as a transmission controller, a slave node 932 operating as an engine sensor and controller, a slave node 933 operating as an air conditioner controller, and a slave node 934 operating as a light indicator controller.

[0091] Similarly, first-tier slave/second-tier master node 924 is connected to another plurality of second-tier slave nodes 941, 942 and 943, each of which may be assigned a sub-sphere of responsibility in the cell controlled by the central section controller. The second-tier slave nodes may therefore include, for example, a slave node 941 operating as a rear door

controller, a slave node 942 operating as a light controller, and a slave node 943 operating as a magnetic breaker controller.

[0092] Each of the first-tier slave nodes 923, 924, 925 and 926 (even if operating as a second-tier master node) may be connected to one or more input/output modules 930. For example, the slave node 925 operating as a front section controller may be connected to a front door control module 951, a kneeling mechanism control module 952, a wheel chair platform control module 953, and a headlight output module 954. Likewise, the slave node 926 operating as a panel controller may be connected to an indicator module 961, an instrument module 962, a control switch module 963, and other miscellaneous modules 964. Virtually any type of input/output or control function may be represented as a module 930. In each instance, the respective slave node 923, 924, 925 and 926 controls the input/output modules 930 connected to it.

[0093] The master node 904 may be connected to a computer 907 through an interface 906 (such as an RS-232 interface), if desired. Through the computer 907, the master node 904 can be instructed to execute certain functions or enter certain control modes. Also, the master node 904 can be monitored or reprogrammed through the computer 907.

[0094] In operation, the master node 904 communicates with the cell controllers 923, 924, 925 and 926 using the main bus 903. The master node 904, as previously described, is generally responsible for delegating control commands to the slave nodes 923, 924, 925 and 926, and for responding to status information and events communicated to the master node 904 over the main bus 903. Each of the slave nodes 923, 924, 925 and 926 receives commands from the master node 904, and issues appropriate commands to their respective second-tier slave nodes 931-934 or 941-943, or input/output modules 930.

[0095] Generally, the slave nodes are disposed in physical locations near the mechanisms which they control. The main data bus 904 and secondary data buses 913, 914 each form a loop connecting the various nodes connected to the bus in a continuous fashion. The data buses 904, 913 and 914 are not restricted to any particular baud rate. Rather, communication may be carried out over each data bus 904, 913 and 914 at a rate that is suitable for the particular control application. Moreover, there is no particular requirement that the data buses in the the Fig. 9

control network (or the more generalized control networks shown in Figs. 7 and 8) be serial data buses. Rather, the data buses may be parallel data buses in situations, for example, where a high data bandwidth is required.

[0096] In the particular control application relating to Fig. 9, each of the nodes is preferably housed in a rugged, potted case made of a suitable lightweight material such as aluminum that provides environmental protection and allows for heat dissipation, as previously described with respect to Fig. 4. In other control environments, other types of housings may be used.

[0097] Fig. 10 is a diagram of a hierarchical control network 1150 utilizing the similar principles as described with respect to the embodiments shown in Figs. 7, 8 and 9, for example, but extrapolated to N tiers. As illustrated in Fig. 10, a first-tier master node 1100 and a plurality of first-tier slave nodes 1102 are connected to the main data bus 1101. Each of the first-tier slave nodes 1102 (up to, e.g., k nodes) may operate as a second-tier master node with respect to a respective second-tier data bus (denoted 11S1, 11S2 ... 11Sk in Fig. 10), each of which has connected to it a plurality of second-tier slave nodes 1103. Likewise, each of second-tier slave nodes 1103 may be connected to a respective third-tier data bus (denoted 11T1, 11T2 ... 11Tk in Fig. 10), each of which has connected to it a plurality of third-tier slave nodes 1104. This same expansion may occur at lower tier levels, with respect to fourth-tier data buses (denoted 11F1, 11F2 ... 11Fk in Fig. 10) down to the Nth level. The control network 1150 depicted in Figure 10 thereby provides a multiple-bus hierarchical control scheme, with each node preferably capable of fanning out to a number of lower-level slave nodes, which may be repeated down to an arbitrary number of levels or tiers.

[0098] In a preferred embodiment, each node in Fig. 10 is given a unique node identifier to facilitate communication among the various nodes in the network. The unique node identifier may comprise, for example, a series of three binary numbers XYZ, the first number X identifying the tier for which the node is connected in its capacity as a slave, the second number Y identifying which bus of the tier (i.e., tier X) the node is attached to by identifying the relative position of the bus with respect to the other buses attached to the tier, and the third number Z

identifying the relative position of the node with respect to other nodes attached to the same bus (i.e., bus Y of tier X).

[0099] While the control network architecture depicted in Fig. 10 has the advantage of being relatively easily expandable and is capable of controlling very large networks, a potential drawback is that the first-tier master node 1100 may have a difficulty in communicating with some of the lowest tier nodes, or an undesirable time lag or latency may occur due to the need to propagate important communications from lower tier nodes through a number of intermediary nodes up to the first-tier master node 1100. Another potential drawback is that a failure of one of the intermediate nodes acting as a master node with respect to slave nodes underneath it is that a failure of such an acting master node may prevent the slave nodes underneath it from communicating with the first-tier master node.

[0100] Accordingly, to overcome one of more of these potential drawbacks, in a preferred embodiment a supervisory network is added to the hierarchical control network and connected to one or more of the lower-tier buses, so as to allow monitoring of communications at the lower tier levels, and to further allow the possibility of redundant backup control and/or communication with respect to the lower-tier master nodes.

[0101] Fig. 11 is a diagram depicting an example of a matrix control network 1290 including a hierarchical control network (generalized to N levels or tiers, as in Fig. 10) with an added supervisory network according to a preferred embodiment as described herein. As illustrated in Fig. 11; the supervisory network comprises a first-tier supervisory master node 1252 connected to a common supervisory data bus 1251. A plurality of supervisory slave nodes 1254 are connected to the supervisory data bus 1251. Each of the supervisory slave nodes 1254 is also connected to one of the lower-tier buses, as illustrated in Fig. 11, and each of the supervisory slave nodes 1254 monitors communications occurring on the particular lower-tier bus to which it is connected. Information monitored by the supervisory slave nodes 1254 is communicated over the supervisory data bus 1251 to the first-tier supervisory master node 1252 either as needed, or on a regular basis. The first-tier supervisory master node 1252 can relay the information to the first-tier master node 1200 over the first-tier common bus 1201.

[0102] In a preferred embodiment, the supervisory master node 1252 and the supervisory slave nodes 1254 each comprise a downlink transceiver and an uplink transceiver, and may be configured, for example, in accordance with the node 603 depicted in Fig. 6. In Fig. 11, the uplink transceiver portion of each node is denoted with an “S” (for Slave), while the downlink transceiver portion of each node is denoted with an “M” (for Master). As illustrated in Fig. 11, the supervisory network may use the same principles and take the same general architecture as the control network shown in Figs. 7 or 8, for example; however, the supervisory network is connected “vertically” so as to connect with one or more lower-tier data buses, thereby allowing interbus monitoring control. The addition of the supervisory network therefore, in one aspect, may be viewed as creating a matrix-type architecture, due to the intersection of control paths emanating downward from a single master node 1200 with the supervisory paths emanating across the control network from a single supervisory master node 1252.

[0103] The supervisory network illustrated in Fig. 11, comprising the supervisory data bus 1251 and associated supervisory master node 1252 and supervisory slave nodes 1254, may in one aspect allow rapid communication between different tiers of the hierarchical control network. In the absence of the supervisory network, information from the lowest tiers can only reach a higher-tier master node (e.g., the first-tier master node 1200) by propagating up each tier through a master/slave data exchange at each level, until eventually it reaches the appropriate higher-tier master node. In the case of an emergency, or where time of response by the higher-tier master node is critical, the delay in communicating the information from the lowest tiers to the higher-tier master node can be detrimental. The supervisory network allows information, especially information that is urgent, to reach the higher-tier master node faster than would be possible in the absence of the supervisory network, by bypassing the required master/slave data exchange at each intervening tier.

[0104] Not all buses or tiers need to be monitored by the supervisory network. However, it is possible that all of the buses in each tier of the main hierarchical control network (i.e., the control network in the absence of the supervisory network) can be connected to a supervisory node, so as to provide comprehensive supervision and monitoring of all of the buses at each level

of the main hierarchical control network. Preferably, only one supervisory node is attached to any given bus in the main hierarchical control network.

[0105] If the control network is large and/or has many levels or buses, then the supervisory network can be expanded by providing additional lower-tier supervisory levels, including lower-tier supervisory buses, configured with a control structure similar to the hierarchical control network of, e.g., Fig. 7, 8 or 10, oriented and connected so as to provide cross-bus communication among the different tiers or buses of the main hierarchical control network. Fig. 12 is a diagram of a matrix control network 1390 with an expanded supervisory network according to a preferred embodiment as described herein, generally capable of monitoring more lower tiers, and in some cases providing more flexible monitoring capability, than the supervisory network shown in the embodiment of Fig. 11. As illustrated in Fig. 12, the supervisory network comprises a first-tier supervisory master node 1352 connected to a first-tier common supervisory data bus 1351. A plurality of first-tier supervisory slave nodes 1354 are connected to the supervisory data bus 1351. Each of the first-tier supervisory slave nodes 1354 may also be connected to its own second-tier common supervisory data bus, such as second-tier supervisory data bus 1361 depicted in Fig. 12, via an intermediary second-tier supervisory master node 1362. One or more second-tier supervisory slave nodes 1364 may be connected to the second-tier supervisory data base 1361.

[0106] Where the fan-out of the first-tier supervisory master node 1352 is limited, the additional second-tier common supervisory data buses may allow for additional supervisory nodes to be connected to lower-tier data buses of the main hierarchical control network, thereby expanding the monitoring capability of the supervisory network. In a preferred embodiment in accordance with the principles and network architecture illustrated in Fig. 12, the first-tier supervisory master node 1352, first-tier supervisory slave nodes 1354, second-tier supervisory master node 1362, and second-tier supervisory slave nodes 1364 each comprise a downlink transceiver and an uplink transceiver, and may be configured, for example, in accordance with the node 603 depicted in Fig. 6. In Fig. 12, the uplink transceiver portion of each node is denoted with an "S" (for Slave), while the downlink transceiver portion of each node is denoted with an "M" (for Master).

[0107] Each of the first-tier supervisory slave nodes 1354 and second-tier supervisory slave nodes 1364 is connected to one of the lower-tier buses, as illustrated in Fig. 12, and each of the first-tier supervisory slave nodes 1354 and second-tier supervisory slave nodes 1364 monitors the communications occurring on the particular lower-tier bus to which it is connected.

5 Information monitored by the first-tier supervisory slave nodes 1354 and second-tier supervisory slave nodes 1364 is communicated, directly or indirectly, over the supervisory data bus 1351 to the first-tier supervisory master node 1352 either on an as-needed or a regular basis. The first-tier supervisory master node 1352 may relay the information to the first-tier master node 1300 over the first-tier common bus 1301.

10 **[0108]** The existence of the second-tier supervisory master node 1362 in the network structure of Fig. 12 allows physical separation of the second-tier supervisory data bus 1361 from the data bus (in this example, second-tier data bus S1) of the main hierarchical control network. The second-tier supervisory master node 1362 is therefore free to communicate with (e.g., poll) the various second-tier supervisory slave nodes 1364, without having to be concerned about
15 potential collisions with data being communicated over the second-tier data bus S1 to the second-tier slave nodes 1311 from the appropriate first-tier slave/second-tier master node 1310. In other words, the second-tier supervisory data bus 1361 is a dedicated bus ordinarily used for supervisory functions, as opposed to control functions.

[0109] The same principle of adding a second-tier supervisory data bus 1361 with
20 second-tier supervisory slave nodes 1364 can be extended by the addition of lower-tier supervisory data buses, just as lower-tier data buses may be added to the main hierarchical control network (as explained earlier herein, for example, with respect to the control network shown in Fig. 10). The result is a supervisory network that is as large and flexible as the situation demands. In addition, speed of monitoring may be increased by providing second or
25 lower tier supervisory data buses. The number of supervisory nodes at a given tier or bus may be reduced by the addition of one or more lower-tier supervisory data buses, thereby reducing the time need to poll the supervisory nodes on the upper-tier data bus. Those skilled in the art will appreciate that the number of tiers and buses in the supervisory network, and the number of supervisory nodes for each tier or bus, will depend upon the particular application and the

structure (i.e., number of tiers and buses, and the arrangement thereof) of the main hierarchical control network.

[0110] As with the network structure of Fig. 11, the resulting network structure depicted in Fig. 12 may be viewed as a “matrix” architecture due to the intersection of control paths emanating downward from a single master node 1300 with the supervisory paths emanating across the control network from a single supervisory master node 1352. It is possible that all of the buses in each tier of the main hierarchical control network can be connected to one of the supervisory nodes, so as to provide comprehensive supervision and monitoring of all of the buses at each level of the main hierarchical control network, and preferably only one supervisory node is attached to any given bus of the main hierarchical control network. It is not necessary, however, that each of the buses of the main hierarchical control network have a supervisory node attached to it.

[0111] While various communication methods may be used in connection with the embodiments shown in Figs. 11 and 12, a polling scheme is preferred over, for example, a voting scheme. In other words, the master node for each bus communicates with each slave node on the bus in a designated time slot. The slave nodes may be polled sequentially, or in some other order, depending upon the particular application. Generally, the slave node waits until the master node sends a polling signal to the slave node before the slave node attempts to communicate over the bus. In this manner, collisions are avoided, as the master node communicates with only one slave node at a given time. Use of a polling communication scheme can be efficient and effective, and does not require the need for arbitration of competing signals, as may be required in a voting scheme. In a voting communication scheme, unlike a polling scheme in which collisions are avoided, a collision of competing signals can require a relatively long time to resolve, and for control to be effectively re-established. Also, the data bits in a data burst can be relatively short in a polling scheme, relative to the size of the data bits used in a voting scheme.

[0112] In a preferred embodiment, communication over a common bus is carried out using time division multiplexing, wherein the master node for a particular bus polls the slave nodes on the bus sequentially or otherwise in different time slots. In a preferred embodiment, a

special supervisory time slot is set aside for communication between the master node and the supervisory node connected to the particular common bus.

[0113] Figs. 13A and 13B are timing diagrams for communication in a hierarchical control network having a matrix architecture, such as illustrated in Fig. 11 or 12, wherein a supervisory time slot is utilized. As illustrated first in Fig. 13A, in normal operation a master node (such as first-tier master node 1200 or 1300, or any other master node) alternates between a “command mode” (or “master mode”) and a “listen mode” (or “slave mode”) within a master node time frame 1400. In a slave node command/listen time segment 1405 of the master node time frame 1400, the master node polls the slave nodes by first commanding and then awaiting a response from each slave node. The master node thus alternates between a slave node command time slot 1411 and a slave node listen time slot 1412, commanding the slave node in the slave node command time slot 1411, and listening to the slave node in the slave node listen time slot 1412. Following the slave node command/listen time segment 1405, the master node time frame includes a supervisory node command/ask time segment 1406 and a supervisory node response time segment 1407. In the supervisory node command/ask time segment 1406, the master node transmits a command to the supervisory node, or else requests certain information from the supervisory node. In the supervisory node response time segment 1407, the master node awaits a response, if any, from the supervisory node.

[0114] The master node time frame 1400 is repeated so as to effectuate control of the hierarchical control network, and also provide for supervisory monitoring functions by the supervisory network. Thus, as illustrated in Fig. 13A, the master node returns to polling the slave nodes (i.e., alternating between slave node command time slots 1411 and slave node listen time slots 1412) at the end of the supervisory node response time segment 1407.

[0115] During the supervisory node response time segment 1407, the master node temporarily enters a listen or “slave” mode. In one embodiment as described herein, if the supervisory node observes an illegal operation by the master node, then the supervisory node can shut the master node down by commanding the master node appropriately during the supervisory node response time segment. Such a situation is illustrated in Fig. 13B, which, like Fig. 13A, illustrates a master node time frame 1450 comprising a slave node command/listen time segment

1455 (including slave node command time slots 1461 and a slave node listen time slots 1462), a supervisory node command/ask time segment 1456, and a supervisory node response time segment 1457. Assuming that the supervisory node detecting, e.g., an illegal operation by the master node during the slave node command/listen time segment 1455, or the supervisory node command/ask time segment 1456, the supervisory node can, during the supervisory node response time segment 1456, instruct the master node to shut down.

[0116] The shut down capability by the supervisory node may be utilized regardless of which tier the master node controls. If the master node is on a tier other than the first tier, then the supervisory node can utilize a secondary means for shutting down the lower-tier master node in case the direct shut down command does not succeed. The supervisory node may send a shut down request to the first-tier supervisory node (such as first-tier supervisory node 1252 in Fig. 12, or 1352 in Fig. 13) via the supervisory network, and the first-tier supervisory node then can transfer the request to the first-tier master node (e.g., first-tier master node 1200 or 1300) which will shut down the lower-tier master node by issuing the appropriate command through the main hierarchical control network channels.

[0117] Generally, in absence of failure, supervisory nodes only listens to the data bus to which it is attached, and does not interfere with the control scheme of the main hierarchical control network. However, in some embodiments, it is possible for the supervisory node to take on additional control functionality.

[0118] In one embodiment as described herein, the first-tier supervisory master node 1252 (or 1352) is programmed to provide redundant backup control for the first-tier master node 1200 (or 1300). In such an embodiment, the first-tier supervisory node 1252 may detect a failure of the first-tier master node 1200 in the same manner as the first-tier slave nodes 1210 do, or by detecting an illegal operation, or both. Upon detecting a failure of the first-tier master node 1200, the first-tier supervisory node 1252 may substitute itself for the first-tier master node 1200, taking over some or all of the functionality of the first-tier master node 1200. Should the first-tier supervisory node 1252 thereafter fail, one of the first-tier slave nodes 1210 can then take over as first-tier master node in place of the first-tier supervisory node 1252, according to the prioritized redundant backup control scheme described earlier herein.

[0119] The above principles may be extended to apply to each of the buses which are monitored by a supervisory node, such that the supervisory node of a given bus can shut down the master node of that particular bus upon detecting a failure of the master node, and thereafter assume the responsibilities and functionality of the master node for that bus.

5 **[0120]** In another embodiment as described herein, the first-tier supervisory master node 1252 (or 1352) is similarly programmed to detect a failure of the first-tier master node 1200 by, e.g., detecting an illegal operation. Upon detecting a failure of the first-tier master node 1200, the first-tier supervisory node 1252 commands the first-tier master node 1200 to shut down, and temporarily takes over some or all of the functionality of the first-tier master node 1200. For
10 example, the first-tier supervisory master node 1252 may temporarily take over only the critical functions of the first-tier master node 1200. Within a pre-programmed failure detection time frame, as previously described herein with respect to redundant backup control, one of the first-tier slave nodes 1210 will detect the fact that the first-tier master node 1200 has failed (as the first-tier master node 1200 has been shut down by the first-tier supervisory master node 1252),
15 and the first-tier slave node 1210 detecting the failure of the first-tier master node 1200 will then automatically take over as the new first-tier master node. When such a substitution occurs, the first-tier supervisory node 1252 relinquishes its responsibilities of the first-tier master node 1200 that it had temporarily assumed, and returns to monitoring the first-tier data bus 1201.

[0121] In such a manner, both effective monitoring of the master node and redundant
20 backup control are provided. The first-tier supervisory master node 1252 need not be programmed with all of the capabilities of the first-tier master node 1200, since one of the first-tier slave nodes 1210 automatically takes over as the first-tier master node within a relatively short period of time.

[0122] Again, the above principles may be extended to apply to each of the buses which
25 are monitored by a supervisory node, such that the supervisory node of a given bus can shut down the master node of that particular bus upon detecting a failure of the master node, and thereafter temporarily assume the responsibilities and functionality of the master node for that bus, until a slave node takes over as the new master node for the bus.

[0123] Other embodiments of network architectures as disclosed herein include one or more multi-master data buses in a supervisory network, for the purpose of monitoring and/or control of the various data buses in the main hierarchical control network. Fig. 14 is a conceptual diagram of a matrix control network 1500 including a multi-master data bus supervisory network 1525 connected to a hierarchical master-slave control network 1501 for the purpose of monitoring and/or supervising the various buses at different level tiers of the hierarchical master-slave control network 1501. As illustrated in Fig. 14, the hierarchical master-slave control network 1501 comprises a plurality of levels or tiers 1503 (1 through N), each tier 1503 conceptually representing one or more data buses at the respective hierarchical level of the master-slave control network 1501. Each of the data buses represented at each tier 1503 is preferably connected to a plurality of nodes organized according to a master-slave hierarchy, such as illustrated, for example, in Fig. 10. The multi-master data bus supervisory network 1525 is coupled by monitoring lines 1505 to the various tiers 1503 of the hierarchical master-slave control network 1501, to allow monitoring of the nodes at the different tiers 1503 of the hierarchical master-slave control network 1501. In one aspect, the multi-master data bus supervisory network 1525 allows rapid communication in the “vertical” direction across the tiers 1503 of the hierarchical master-slave control network 1501, thereby allowing more rapid response by the top-tier master node than might otherwise be possible to events needing urgent attention, such as exception or error conditions.

[0124] Details of a preferred matrix control network 1690 with a multi-master data bus supervisory network are illustrated in Fig. 15. The primary hierarchical control network shown in Fig. 15 is similar to that shown by itself in Fig. 10, and in connection with other supervisory network structures in Figs. 11 and 12. In addition to the primary hierarchical control network, a supervisory network is added and connected to one or more of the control network data buses, so as to allow monitoring of communications at the mid- and lower-tier levels, and to further allow the possibility of redundant backup control and/or communication with respect to the master nodes at the various different levels.

[0125] In Fig. 15, the matrix control network 1690 (generalized to N levels or tiers, as in Fig. 10, 11 or 12) comprises a main hierarchical control network including a first-tier master

node 1600 connected to a first-tier data bus 1601 (designated F1 in Fig. 15), and a plurality of first-tier slave nodes 1610 connected to the first-tier data bus 1601. Each of the first-tier slave nodes 1610 is capable of acting as a second-tier master node with respect to a second-tier data bus (designated as S1, S2 ... S(k-1) in Fig. 15), and each of the second-tier data buses preferably has a plurality of second-tier slave nodes 1611 connected to it. The main hierarchical control network may be expanded downward in a similar manner for an arbitrary number N levels, with each of the second-tier slave nodes 1611 potentially being connected as third-tier master nodes to third-tier data buses, and so on. The main hierarchical control network may have all of the functionality and features of the control network 1150 illustrated in Fig. 10, for example, and described earlier herein.

[0126] As further illustrated in Fig. 15, a supervisory network is preferably added to the main hierarchical control network for the purposes previously described herein. Accordingly, a first-tier multi-master common supervisory data bus 1651 is provided in the matrix control network 1690, and a plurality of first-tier supervisory nodes 1652 and 1660 are connected to the first-tier multi-master common supervisory data bus 1651, including a “top-level” first-tier supervisory node 1652 connected to the first tier data bus 1601 and a plurality of other first-tier supervisory nodes 1660 each connected to a lower-tier data bus of the main hierarchical control network. Each first-tier supervisory node 1652, 1660 preferably comprises an uplink transceiver and a downlink transceiver, and may generally be configured, for example, in accordance with the node 603 depicted in Fig. 6. In Fig. 15, the uplink transceiver portion of each node is denoted with “S” (for Slave), while the downlink transceiver portion of each node is denoted with “MM” (for Multi-Master).

[0127] In operation, the top-level first-tier supervisory node 1652 monitors activity on the first-tier data bus 1601, and each of the other first-tier supervisory nodes 1660 monitors activity on the lower-tier data bus of the hierarchical control network to which it is connected. If an event needing urgent attention occurs on one of the lower-tier data buses, then the first-tier supervisory node 1660 detecting the event sends a message over the first-tier multi-master common supervisory data bus 1651 to the appropriate supervisory node, typically the top-level

first-tier supervisory node 1652 (although the error message may be broadcast to all of the first-tier supervisory nodes 1652, 1660).

[0128] The top-level first-tier supervisory node 1652 is polled by the first-tier master node 1600 periodically, as if it were any of the first-tier slave nodes 1610. When the first-tier master node 1600 polls the top-level first-tier supervisory node 1652, the first-tier master node 1600 learns of the event requiring urgent attention, and responds accordingly by either issuing commands down the hierarchical control network, or else optionally, if only a limited response is needed, sending a responsive message to the lower node or data bus through the first-tier supervisory nodes 1652, 1660 via the first-tier multi-master common supervisory data bus 1651.

When initially polling the top-level first-tier supervisory node 1652, the first-tier master node 1600 generally limits its inquiry to whether or not an error or exception has occurred at some lower level. If the top-level first-tier supervisory node 1652 answers the poll by indicating an error or exception has occurred, then the first-tier master node 1600 sends a control message to the top-level first-tier supervisory node 1652 requesting specific details of the exception or error condition. Using this technique, the first-tier master node 1600 is able to rapidly learn of events requiring urgent attention at even the lowest tiers of the hierarchical control network, and to respond quickly thereto.

[0129] The first-tier supervisory nodes 1652, 1660 connected to the first-tier multi-master common supervisory data bus 1651 may use a conventional protocol to gain access to and communicate over the first-tier multi-master common supervisory data bus 1651. Conventional protocols (including Ethernet protocols) exist in the art of local area networks that may be used or readily adapted in the context of the first-tier multi-master common supervisory data bus 1651. Such protocols generally include mechanisms for handling collisions between competing nodes attempting to gain control of the common bus at the same time. For example, if a cyclic redundancy check (CRC) field for the received message does not match the expected value, then a collision or other error is assumed, and the competing supervisory nodes will access the multi-master common supervisory data bus 1651 bus according to a relative, predetermined priority. However, collisions are expected to be very rare, since communication over the first-tier multi-master common supervisory data bus 1651 is only necessary when an exception condition

requiring attention is detected at one of the lower tiers. As a result, the main hierarchical control network maintains a very efficient master-slave communication protocol (i.e., a polling protocol) for allowing the master nodes at the various tiers to control the slave nodes at each such tier, while at the same time providing a mechanism for intermittent rapid communication between the lower tiers and upper or first tier via supervisory nodes relying upon a different protocol suitable for a network having multiple master nodes but only requiring occasional communication between the master nodes.

[0130] The supervisory network may comprise additional multi-master supervisory data buses connected to the first-tier data bus 1601 through one or more supervisory nodes (e.g., first-tier supervisory node 1653), to allow monitoring of additional data buses of the hierarchical control network. Similarly, the supervisory network may include one or more second-tier supervisory data buses, such as second-tier multi-master supervisory data bus 1661, to provide additional supervisory nodes for monitoring data buses of the hierarchical control network. As illustrated in Fig. 15, the second-tier multi-master supervisory data bus 1661 may be connected to a plurality of second-tier supervisory nodes 1672, 1680, including a top-level second-tier supervisory node 1672 and a plurality of other second-tier supervisory nodes 1680. As with the first-tier supervisory nodes 1652, 1660, each second-tier supervisory node 1672, 1680 preferably comprises an uplink transceiver and (denoted “S” in Fig. 15) a downlink transceiver (denoted “MM” in Fig. 15), and may generally be configured, for example, in accordance with the node 603 depicted in Fig. 6. In operation, the top-level second-tier supervisory node 1672 monitors activity on the second-tier data bus 1602, and each of the other first-tier supervisory nodes 1680 monitors activity on the lower-tier data bus of the hierarchical control network to which it is connected. If an event needing urgent attention occurs on one of the lower-tier data buses, then the second-tier supervisory node 1680 detecting the event sends a message over the second-tier multi-master common supervisory data bus 1661 to the appropriate supervisory node, typically the top-level second-tier supervisory node 1672. The top-level second-tier supervisory node 1672 is polled by the appropriate second-tier master node 1610 periodically, as if it were any of the second-tier slave nodes 1611. When the second-tier master node 1610 polls the top-level second-tier supervisory node 1672, the second-tier master node 1610 learns of the event requiring

urgent attention, and responds accordingly by either issuing commands down the hierarchical control network, or else, if only a limited response is needed, sending a responsive message to the lower node or data bus through the second-tier supervisory nodes 1672, 1680 via the second-tier multi-master common supervisory data bus 1671, or else relaying the message to the first-tier master node 1600 when polled by the first-tier master node 1600, or any combination of these responses. Also, the first-tier master node 1660 (if any) connected to the second tier data bus 1602 may also learn of the event by listening to communications across the second tier data bus 1602, and may respond or else relay the information to the first-tier master node 1600, if desired.

[0131] Using the above techniques, the first-tier master node 1600 and other master nodes (such as first-tier master node 1610) are able to rapidly learn of events requiring urgent attention at even the lowest tiers of the hierarchical control network, and to respond quickly thereto.

[0132] The second-tier supervisory nodes 1672, 1680, as with the first-tier supervisory nodes 1652, 1660, may use a conventional protocol to gain access to and communicate over the second-tier multi-master common supervisory data bus 1661. As before, collisions are expected to be infrequent, since communication over the second-tier multi-master common supervisory data bus 1661 is only necessary when an urgent condition requiring attention is detected at one of the lower tiers.

[0133] In one aspect, the addition of the supervisory network to the main hierarchical control network may be viewed as creating a matrix-type architecture, due to the intersection of control paths emanating downward from a single master node 1600, with supervisory paths emanating across the control network from the various supervisory nodes 1660. In the absence of the supervisory network, information from the lowest tiers can only reach a higher-tier master node (e.g., the first-tier master node 1600) by propagating up each tier through a master/slave data exchange at each level, until eventually it reaches the appropriate higher-tier master node. In the case of an emergency, or where time of response by the higher-tier master node is critical, the delay in communicating the information from the lowest tiers to the higher-tier master node can be detrimental. The supervisory network allows information, especially information that is urgent, to reach the higher-tier master node faster than would be possible in the absence of the

supervisory network, by bypassing the required master/slave data exchange at each intervening tier.

[0134] Not all buses or tiers need to be monitored by the supervisory network. However, it is possible that all of the buses in each tier of the main hierarchical control network (i.e., the control network in the absence of the supervisory network) can be connected to a supervisory node, so as to provide comprehensive supervision and monitoring of all of the buses at each level of the main hierarchical control network.

[0135] In one aspect, the supervisory network may allow faster response by the first-tier master node or other upper-tier master nodes to situations requiring attention in the lower tiers.

An exemplary application of the control network of Fig. 12 or 13 is a building control system which electronically controls the various systems and sub-systems of a building (e.g., air conditioning, fire alarm, power, and so on). In such an exemplary embodiment, the first-tier master node may be, for example, the top-level building controller. The second-tier master nodes may each be building wing controllers. The third-tier master nodes may each be floor controllers for their respective wings. The fourth-tier master nodes may each be office controllers for their respective floor and wing. The fifth-tier master nodes (or fourth-tier slave nodes) may each be individual room controllers for their respective floor, wing and office. A supervisory network may comprise supervisory nodes connected to each of the lowest-tier buses. Should, for example, a fire or other emergency occur in one of the rooms of the building, the supervisory network could immediately relay the information to the top-level building controller (i.e., the first-tier master node), which could take the appropriate action, without requiring the delay of intermediate communications from tier to tier until the top level tier is reached. Furthermore, the supervisory node first identifying the emergency or exception condition may broadcast it to all of the other supervisory nodes on the supervisory data bus, thereby allowing the knowledge to be quickly propagated throughout the entire system.

[0136] In another exemplary application, the hierarchical control system including the supervisory network may be used to control a robotics system. The hierarchical control system can readily be miniaturized and is easily expandable, making it ideal for robotics applications.

[0137] While preferred embodiments are disclosed herein, many variations are possible which remain within the concept and scope of the invention. Such variations would become clear to one of ordinary skill in the art after inspection of the specification and drawings herein. The invention therefore is not to be restricted except within the spirit and scope of any appended
5 claims.